

Analogue-digital systems with modes of physical behaviour

Edwin Beggs¹ and John V. Tucker²

¹ Department of Mathematics, College of Science,
Swansea University, Singleton Park, Swansea, SA2 8PP,
Wales, United Kingdom
`e.j.beggs@swansea.ac.uk`

² Department of Computer Science, College of Science,
Swansea University, Singleton Park, Swansea, SA2 8PP,
Wales, United Kingdom
`j.v.tucker@swansea.ac.uk`

Abstract. Complex environments, processes and systems may exhibit several distinct modes of physical behaviour or operation. Thus, for example, in their design, a *set* of mathematical models may be needed, each model having its own domain of application and representing a particular mode of behaviour or operation of physical reality. The models may be of disparate kinds – discrete or continuous in data, time and space. Furthermore, some physical modes may not have a reliable model. Physical measurements determine modes of operation. We explore the question: *What is a mode of behaviour? How do we specify algorithms and software that monitor or govern a complex physical situation with many modes? How do we specify a portfolio of modes, and the computational problem of transitioning from using one mode to another mode as physical modes change?* We propose a general definition of an *analogue-digital system with modes*. We show how any diverse set of modes – with or without models – can be bound together, and how the transitions between modes can be determined, by constructing a topological data type based upon a simplicial complex. We illustrate the ideas of physical modes and our theory by reflecting on simple examples, including driverless racing cars.

No plan survives contact with the enemy.

After Helmuth von Moltke the Elder (1800-1891)

There are known knowns; there are things we know that we know. There are known unknowns; that is to say, there are things that we now know we don't know. But there are also unknown unknowns – there are things we do not know we don't know.

Donald Rumsfeld (2002)

1 Introduction

A typical analogue-digital system is a system in which a continuous physical environment, process or component is monitored or governed by a discrete algorithmic process. This simple description covers an astonishingly large range of systems, for example: classical control systems for machines, industrial plant, and buildings; networked products and vehicles; human monitoring and surveillance systems; and scientific experiments, mediated by software, that measure physical quantities. Many applications involve hybrid systems and cyber-physical systems³ These are types of analogue-digital systems whose deployment is vast in scope and whose formal theories and design methods have matured; in particular, hybrid systems have a solid theoretical basis [25], which we will discuss shortly (in subsection 1.3). We will take a fresh look at such systems and raise and offer answers to some foundational questions for a general theory of analogue-digital systems, reflecting on the interface between the physical and the algorithmic, and defining formally a new class of complex analogue-digital systems having distinct modes of physical behaviour.

1.1 Observations on real-world analogue-digital systems

To begin, consider from first principles, a real-world analogue-digital system that consists of physical equipment that is controlled by software on a processor. It is analogue-digital because the system involves continuous and discrete data. Typically, the data characterising the equipment is represented by real numbers, and the data processed by the program is represented by bits. The equipment is made from physical components that exchange data with the controlling software. An important design criterion is the

Principle 0: Robustness. *The design of the software and its underlying algorithms must be able to cope with any data that may be output from, or input to, the physical components.*

Typically, the design of the software is based upon

- (a) mathematical models – generally called *dynamical systems* – of the behaviour of the physical equipment;
- (b) exception handlers if none of the mathematical models apply to a mode of behaviour;
- (c) logical and algebraic models of the behaviour of the programs.

Of course, in the case of (c), the behaviour of the programs is derived from the semantics of the specification and programming languages employed; we will not consider this factor. Consider (a): how the software depends on models of the behaviour of the equipment by, or derived from, dynamical systems.

³ A cyber-physical system generalises the concept of an embedded system to a network of interacting devices with physical input and output [23]. For example, cyber-physical systems and the internet of things are central to speculations on the future of manufacturing in [15].

To create the software for an analogue-digital system, the physical must be replaced by an abstract specification that documents certain operations, tests and properties and that constitutes a data type interface between physical quantities and the algorithms and software. The interface must enable portability and verification so that the software can be certified “as reliable as possible” according to some best practice engineering standard. For the designer of an analogue-digital system, specification and validation presents certain problems outside software engineering. The reliability of an analogue-digital system depends upon the abstract assumptions about the physical system, the sensors and actuators. These assumptions are based on dynamical systems that model physical reality.

A dynamical system is a mathematical model of an entity, process, or environment whose behaviour changes in time. The model represents behaviour by means of states that change over time. The model is likely to have parameters that are external inputs, which represent either

- (i) external influences that are not predictable physical factors to which the system must respond or, conversely,
- (ii) known parameters or instructions designed to control the system.

The state of a dynamical system can be based on continuous data, discrete data or a combination of both.

Complex environments, processes and systems may exhibit quite *distinct* modes of physical behaviour or operation. Thus, for the design of complex physical environments, processes and systems, a *set* of mathematical models may be needed, each model having its own physical domain of application, and representing a particular mode of behaviour or operation of physical reality. The models may be of disparate kinds. Certainly, no single mathematical model is adequate. Furthermore, the models may not cover adequately all the possible modes of behaviour, i.e., not all physical modes may have a reliable model.

1.2 Modelling analogue-digital systems with modes

This paper is about complex analogue-digital systems with many modes. We address the following questions: In a complex analogue-digital system,

What is a mode of behaviour?

How can modes of physical behaviour, with a portfolio of different models, be specified for the use of algorithms and software?

How can modes of physical behaviour, without sound models, be specified for the use of algorithms and software?

How are all the modes bound or linked together to make a robust system? As the system operates, one model becomes less relevant and another more relevant model must be chosen to replace it. What conditions govern the transition from one mode to another?

We introduce initial concept of an analogue-digital system of the form
physical system + interface and protocol for data exchange + algorithm.

Specifically, we propose the algorithms and software treat the physical system as an *oracle*, whose queries are mediated by a protocol governing an interface.⁴ A good deal is known about the basic properties of such a physical oracle model in cases where the physical systems are *very* simple and can be faithfully captured by a *single* dynamical system and, therefore, have one and only one mode [6,7,8,9,10,12]. Here we generalise this concept of an analogue-digital system by introducing modes that may have models or need to be handled as exceptional situations. See the schematic in Figure 1 with four modes.




Figure 1: Illustrating four modes

We propose a theoretical framework for the development of such analogue-digital systems with modes of physical behaviour and their evaluates their fitness. Essentially, we define the concept of a mode to be a data type representing a physical behaviour and equipped with criteria that evaluate its accuracy or relevance.

To the interface and data exchange protocol are added a set of *mode transition functions* that specify what needs to be done to make a transition from one mode to another. We create a simple geometric description of the system, one which describes both the relationship between the various modes of operation of the system and when transitions between the modes are necessary. Such a geometry needs a topological space and so we construct a simplicial complex – a space made of points, lines, triangles, tetrahedra etc., glued together using simple rules – that we call a *nerve*. The points represent the modes, and the lines or triangles joining them are used to show how well each mode describes the physical system. As this description is very explicit, it forms a data type that can be used represent visually the mode transition protocol in the design of the algorithmic structure of the portfolio. Working from first principles, we offer a new perspective that is distinct from hybrid systems.

⁴ The idea is a generalisation of Turing’s and Post’s ideas about oracles to algorithms in computability theory.

1.3 Hybrid systems

Over the years, the term hybrid system has had a number of meanings, such as combining analogue and digital computation [38]. Currently, the term usually refers to a *control* system with analogue and digital behaviour [25], which is an interpretation suitable for our investigation here. Yet, as the ways of studying such systems are many and growing there is a reluctance to define the term (cf. [22]). We prefer to speak informally of analogue-digital systems for which we are also studying general computability theories [1,11,39,40]. The computational aspects of designing hybrid systems received a great impetus in the 1990s from studies by Zohar Manna, Amir Pnueli and Tom Henzinger [26,17], who developed formal methods for modelling, based on finite automata and differential equations, and for reasoning, based on the use of modal and temporal logics and model checking. This has been a foundation for extensive research, tool building and applications in computer science. Thus, for some the term hybrid systems suggests work on control systems using these formal methods. We will summarise some features of the formal methods for modelling to help clarify the ways our approach is distinct; we return to these points in the concluding remarks (11.1).

Hybrid systems are studied by petri nets or hybrid automata. In their graphical representation, the boxes in the hybrid automata are called *locations* or *finite states*.⁵ The boxes can have *invariants* - conditions which must be satisfied as long as the system is in that state, and if they stop being satisfied, then the system must move to another state. The arrows between the boxes represent the allowed state transitions and have *guard conditions*, which must be satisfied for the system to move along that arrow. The arrows can also have associated *actions*, which specify instantaneous (or nearly so) changes of variables on taking the arrow. The states typically have equations for continuous changes of variables, e.g. differential equations. In some implementations we can have hierarchies of finite states, where opening the box for one finite state reveals another hybrid automaton inside: Ptolemy [24] takes such a hierarchical approach with high and low levels of description (the top level components are referred to as modes).

To contrast this with what we are proposing, our key term mode is intended to categorise and modularise physical behaviours. Crucially, the modes of behaviour overlap which means there are transitions between modes. Theoretically, the modes are sets of states of the system; they may have mathematical models but *a priori* they are not associated with any software specification. For each mode α we have a set of data X_α necessary to describe the system, rules for its evolution and (subtle) notions of modes overlapping. To tackle the transition from a set of modes we construct the *nerve*, which is simply the combinatorial representation of the geometry of the system, based on considering the overlapping of the subsets defining the modes.

A hybrid automaton may be nondeterministic, meaning that more than one arrow from a state may have its guard condition satisfied. For our modes, how-

⁵ Terminology varies in the literature – sometimes the finite states in the automaton are called modes.

ever, the structure of the nerve and the idea that all information has errors forces nondeterminism. In fact, where three modes have a simultaneous intersection, we expect a whole continuum of possibilities ranging between two arrows each being an ‘obvious’ one to take. This nondeterminism is linked to the geometry of the state space. From a topological point of view, it corresponds to homotopic paths through state space. The continuum of possibilities in evaluating the arrows is encapsulated into a *mode evaluation* function, which lies at the heart of our approach.

It is instructive to ask a question on making a change of discrete state allowed by guard clause, or in our terms a mode transition. Is this change determined by the behaviour of the system, or is it within the software’s control (ultimately the software designer’s control) to decide what to do? In terms of the popular heating system example, is it the laws of physics which dictate that the system changes, or does a user decide that the room is too hot? We separate out these possibilities to clarify our thinking and to emphasise that it is really the real world system which is in control – the software can only respond to data about the physical behaviour; the distinction may enhance the reusability of the software. Every mode comes with an externally imposed set of orders. The mode evaluation function, which amongst other jobs evaluates how well the current mode is performing, does so by reference to the orders. Some discrete transitions are forced by the system, and others are chosen with reference to the orders.

Finally, we should refer to the Zeno condition, which is that the system should not try to execute infinitely many transitions in finite time. From our point of view, this can’t happen, as the transitions have definite finite durations – in fact estimates of the transition times would be required. To explain why this works, part of the job of mode evaluation is to estimate whether a potential mode to transfer to has a long enough expected duration to evaluate its own performance (including real world measurements) and its possible transfers out – do not transfer into a mode unless that mode can be transferred out of in an orderly fashion. Of course, because of errors in measurements, unexpected behaviour of the system, or just bad programming, that may not be possible, and we may be left with various sub-optimal choices, a matter which we shall discuss – that is our version of the Zeno problem.

1.4 Structure of the paper

The structure of the paper is this. In Section 2, we formulate some general principles for analogue-digital systems to guide our thinking and mathematical formulations. In Section 3, we attempt to clarify some intuitions about modes of operation, and transitions between modes, by introducing an example of a racing car under the control of software.

In Section 4, we give the general definition of an analogue-digital system with modes, and introduce a simple software architecture to organise our discussion. The next tasks are to show how to formalise modes and mode transitions, i.e.,

- (i) to define modes;
- (ii) to evaluate their fitness for purpose; and

(iii) to change from one mode to another.

In Sections 5 we address (i) and in Section 6 we address (ii). In Section 7, the transition theory is developed to address (iii).

Next, we turn to formal examples to illustrate the theory. In Section 8, we give a general mathematical example that characterises many classical systems based on manifolds and geodesics. These illustrate modes in the cases where one can usefully imagine the presence of some global state space based upon combining the state spaces of many modes. In Section 9 we describe an example based on the informal car example in Section 3. In Section 10, we describe an example based on the orbits of planets.

Finally, in Section 11, we return to hybrid systems and touch on topics arising from the study: networks of analogue-digital systems; multi-scale analysis, when we look at the system in finer or coarser detail; and security.

We thank Felix Costa (Lisbon) for many enjoyable and influential conversations on analogue-digital systems; this paper has grown out of our collaboration on [7,6,8,9,10,12].

2 Some principles for analogue-digital system specification

We propose to make a theory that can speak about three distinct components of a complex analogue-digital system: (i) the physical system as it is in reality; (ii) the modes of behaviour of the physical system; and (iii) the algorithms and software that employ the modes to control the physical system. Let us clarify these three distinctions by adopting five more working principles.

Principle 1: Physical uniqueness. *A physical system, process or environment is based on equipment made from particular physical components and whose behaviour is a continuous physical phenomenon that is unique to that equipment in space and time.*

Principle 2: Physical observation. *A physical system, process or environment does what it does. To observe and measure a system is to abstract aspects of its behaviour in a way that is specific to the system. The observations and measurements are the basis of controlling the physical situation. They are the basis for interfacing algorithms and physical behaviour, and for classifying the behaviour of the system into modes.*

Principle 3: Physical specification. *To model a mode of behaviour of a system by a dynamical system is to specify an aspect of the physical system in an abstract and general way; the model specifies a class of particular physical systems having that mode of behaviour. In addition to providing understanding of how the physical system might behave, dynamical systems provide a specification of a general data type interface for a mode.*

The dynamical systems are central to the study of the software. The models of a physical system may be many and varied; they are shaped by different choices of physical insights, spatial and temporal scales, and computational constraints, etc. There are many ways of modelling analytically (such as partial differential equations, ordinary differential equations) and computationally (such as finite elements, neural networks, lattices, cellular automata.)

Principle 4: Modes and Models. *A complex system will have many modes. Ideally, each mode will have a sound model. A portfolio of modes and models, with overlapping domains of application, is needed to cover the system's behaviour.*

Even if each mode of behaviour has its own individual model with a state space, the models can be disparate, and the system does *not* have a unified model or state space.

Recalling the *Principle 0: Robustness* from the Introduction, the aspiration is that all physical eventualities are covered:

Principle 5: Mode Transitions. *A methodology for managing a portfolio of models is needed that has methods for*
(i) evaluating the quality of the models in the portfolio;
(ii) passing from one model to another; and
(iii) dealing with exceptional situations when no model in the portfolio applies.

The fact that the modes are physical processes observed by measurements that are liable to errors is fundamental to our investigation.

3 Physical modes and why they may change

We begin by exploring the need for modes of operation and transitions between modes, informally, by means of a thought experiment on autonomous cars.

3.1 Example of a system with many modes: a robotic racing car

Consider an autonomous car racing championship. To create the software to govern the racing car, the programmers are given:

Observables. Specifications of the data available from the sensors and actuators, and of their behaviour; for example, “the maximum acceleration is $10m/sec^2$ ”; “the time delay on the speedometer reading is $0.5sec$ ”; ...).

Objectives. Driving tactics for the race; for example, “minimise the lap time”; “don’t get too close to the other cars or the edge of the track”; “beware reckless car A”; ...).

First, the programmers must separate concerns to simplify the problem. Immediately, the behaviour of the physical system is split into modes, chosen to simplify the mathematical modelling of observables of the behaviour and the specification of driving objectives. For our racing car, we might manage with five modes as follows:

Mode α : The car is on the track and some distance from the nearest cars.

The state of any other car is reduced simply to the position of a point on the track, as it is not particularly relevant. The important objectives are maximising acceleration and taking corners at the fastest safe speed; and perhaps some longer term tactics, such as when to refuel or do a wheel change.

It would be nice to stay in mode α all the time, but this is unlikely as at some point the car will meet rival cars. This meeting can be determined from the picture available from the state space of mode α by looking at the distance to the nearest cars. When this distance becomes smaller than a certain value, the car will need to change behaviour. This change in separation ought to be expected and predicted by a model; but it could also be unexpected (e.g., the car hits a pool of oil and skids).

Mode β : The car is on the track but close to other cars.

The objectives are to avoid collision, and to overtake if possible. Now more state variables are needed to handle the nearby cars. For example, to positions of the other cars, their velocities, accelerations, and sizes may be added. Also, perhaps, their identities and some intelligence may be added (e.g., “car A always tried to block us when we tried to overtake, car B did not”).

At a particular combination of positions and velocities of nearby cars, a collision may be flagged up as a real danger.

Mode γ : The car is on the track but on a crash course with another car.

The objective becomes not only to see if it is possible to avoid a collision, but to mitigate the effects if it cannot be avoided. At this point the actual size and shape of the cars may be important, if (say) a wheel on wheel impact was likely to be more dangerous than a side on side impact.

These are three simple modes that make an initial design of a system; maybe three modes could win a race, but there are gaps. Suppose the car comes to a complete stop ending up off the track, undamaged on a grass verge: the programmers need to bolt on another mode:

Mode δ : The car is off the track.

The objective is to rejoin the track as soon as possible, but to stay away from cars on the track.

Experience generates further problems that must be considered. Suppose the instrumentation develops a fault and core data is untrustworthy. As the software was not designed to cope with (say) incorrect readings from the speedometer, a crash could result. To fix this the programmers need to add another mode:

Mode ϵ : The car is on the track but the instrumentation is faulty.

The objective is to slow down, avoid other cars and to make way to the grass verge and stop. To add this mode, the programmers need to carry out some modifications to the previous ones. These modifications are not major, they just need to say when the error mode ϵ should be entered. For example, a comparison of the speedometer reading can be made with the changes of positions in time of objects on the side of the track. If there is an obvious disagreement, go to mode ϵ . This modification is for modes α and β : if we are in mode γ , we already have more urgent problems.

A useful distinction is to declare the modes to be *normal* or *exceptional*, for example: modes α , β , and γ are normal and δ and ϵ are exceptional.

3.2 A modular approach to modelling complicated physical systems

From the point of view of designing an analogue-digital system, such as a driverless racing car, the modes represent a form of physical modularisation, determined by the science of the physical domain and, possibly, exceptional circumstances. Consider two aspects of physical modelling: modularity and accuracy.

Modularity. The benefits of modularity are well known, but here modularisation is a necessity. Physical systems possess all sorts of modes, not least modes for startup, normal running, exceptional situations, and shutdown. Mathematical models of physical behaviour are approximations with limited ranges of validity. In complex environments and systems, physical conditions change, and the range of validity of a model can easily be stressed and broken. Thus, it is to be *expected* that *a change of physical conditions requires a change of mathematical model*. In this way complex systems acquire different modes, where the entire description of the system may change, and a completely new set of state variables and a new type of model is used.

The modes constitute a modular approach to understanding physically the system so that

Specification. We can track of properties that characterise a mode and trigger need for a change.

Implementation. When a mode is modified in a minor way, or more special cases added, we can isolate and limit the changes – and consequent errors – in software.

Efficiency. We can detect if a problem can be solved faster with parallel processing.

The models for the modes may be of disparate kinds: they may be analytical models based on PDEs or coupled systems of ODEs; or they may be algorithmic models based on numerical approximation methods, neural networks, coupled map lattices, and cellular automata. In a complex system one can expect several kinds to be present.

There are methods of combining algorithmic models in different domains because algorithmic models are discrete space, discrete time models and can be unified in a theory of synchronous concurrent algorithms [37]; for example, in the case of whole heart modelling, see [28,29,14].

Accuracy. In a simple physical system, with easy access to measurements, it may be reasonable to assume that the software has instant access to a comprehensive specification of the physical system at any time. However, in practice, the system may be complicated enough to require some factors to be estimated. It would be helpful to have a modular approach to modelling the system that can cope with factors such as:

Delay. Measurements take time to perform, or be available only to a limited accuracy. There may be a time-lag in obtaining some information that requires action by the software.

Decay. Sensors, actuators and communication links are all prone to degradation (or even complete failure) and replacement will take time, if it is possible at all. If the performance of a component has degraded, the software may have to spot this from inconsistencies in the data, and take appropriate action.⁶

History. If there is no recent data to base a decision on, historical data may be needed to make a best guess.

4 Analogue-digital systems

We define a general analogue-digital system without modes and then generalise the definition by introducing modes.

4.1 What is a general analogue-digital system?

The basic architecture of an analogue-digital system is depicted in Figure 2, in which arrows are used to indicate the direction of flow of information.




Figure 2: A analogue-digital system

⁶ The consequences of unreliable approximation to reality can be catastrophic, for example, on an airliner crash in the Atlantic following bad weather:

Temporary inconsistency between the measured airspeeds, likely following the obstruction of the Pitot probes by ice crystals that led in particular to autopilot disconnection and a reconfiguration to alternate law, – Final Report On the accident on 1st June 2009, Bureau d’Enquêtes et d’Analyses pour la sécurité de l’aviation civile

The *physical system* is a process, system or environment that we wish to monitor or control. Suppose there is a fixed system of sensors for measurements, and actuators for control, linked to the physical system; so we should think of information flow along a *sensor bus* (e.g., from thermometers) and an *actuator bus* (e.g., to motors).

The *interface* is what translates between the sensors and actuators and the decision making software. The interface can take many forms; in a very simple case of the direct control of a motor by a microprocessor, it would be little more than some digital-to-analogue (DA) and analogue-to-digital (AD) converters. The interface may need to manage error margins and timing delays.

A *query* is a request from the decision making software to the interface and a *response* is a message going the other way. In the motor and microprocessor example, the query is simply a binary number interpreted by the interface's DA-converter as a voltage applied to a motor, and there is no response. The response to a query normally takes time and it is possible that the response is timed out.

Finally, there is the *algorithm* that formalises the decision making that attempts to monitor or control the physical system. The data must be considered separately, as it is the sole basis by which the algorithm knows what is happening in the physical system, it forms the algorithm's picture of physical reality. The data must be updated, either periodically by standing orders in the interface, or when the algorithm sends a query to update items.

The remaining components of Figure 2, mentioned in italics, are the basis of *specifications*. The *laws or theories* refer to mathematical models of the physical system under consideration, e.g., Newton's laws of motion or characteristics of electronic devices. Using these laws, and the sensors and actuators, an *axiomatic specification* of the interface is written.⁷

The axiomatic description of the interface is used to write an algorithm to satisfy the list of *objectives*.⁸ For given initial configurations, if it can be proved that the algorithm suffices to guarantee the objectives then the algorithm can be said to be *verified assuming the validity of the axioms*.

4.2 Physical Oracles

The central idea is that *the physical system or environment is an oracle to the algorithm*. The algorithm requests and receives data about the physical environment or system to use in its processing. The interface between the analogue and digital is of central theoretical interest. Thus, we have in mind a *physical oracle*, in which the computer may *query* the interface, and receive messages back. The queries may be of various forms, e.g., a question for a temperature sensor, or an

⁷ For example, such axioms might be:

“System variable x can be measured to accuracy ϵ in time t .”

“System variable y obeys the differential equation to within error ϵ .”

“System variable z always remains in the interval $[-\pi, +\pi]$.”

⁸ For example, such an objective might be: “System variable x should always satisfy $x \leq 30$.”

instruction for a motor to be turned on. The algorithm need not know anything about the detailed operation of sensors and actuators, that is all abstracted in the interface.⁹

The idea of an oracle is very general and is not confined to computing systems. For example, a complicated socio-technical example would be a laboratory service in a hospital. Instructions are received from doctors to perform tests, and the results are sent back to the doctors. There is a division of labour: the laboratory staff do not need to know about diagnosing disease, and the doctor does not need to know how to perform the tests. The different times taken for tests, the queue of requests and the capacity of the lab mean that the lab staff have to make scheduling decisions. If a request is urgent, they may have to decide between a quick but less reliable method and a slower more reliable test, of course notifying the doctor of the possible errors concerned. In this example, a query by a doctor might be “measure levels of X in patient Y by 09.00 tomorrow”, and a response might be “level is Z, less accurate method used due to urgency of test”. The choice of test is a choice of mode. The idea of oracle can be applied similarly to any service in which time is a prominent factor.

4.3 An analogue-digital system with modes of operation

Consider an analogue-digital system whose physical behaviour can be split into *modes*. Starting with Figure 2, we imagine an architecture for three modes α, β, γ depicted in Figure 3:




Figure 3: A three mode analogue-digital system

Domain scientists and engineers must classify the behaviour of the physical situation to (i) define the modes and (ii) create the list of objectives for each mode. We suppose each mode $\alpha \in \mathcal{M}$ has a list of state variables, a state space and trajectories that represent behaviour in time. These may be derived from a model that is a reasonable mathematical description of the physical system in

⁹ The interfaces are far more complicated than the traditional oracles of computability theory. There, since Emil Post’s development of computability theory, a query asks if a datum is in a set, and the response is yes or no and takes one time step.

that mode, or from some ad hoc rules intended to cope with unusual situations. In each mode, appropriate approximations may be made to keep the description of the mode manageable.

Thus, in the analogue-digital system a component state_α managing the system in mode α will contain the state variables for mode α , and probably more, such as estimates of the errors of the state variables, and some of the known properties of the system.

In addition to coding the modes $\alpha \in \mathcal{M}$, there must also be conditions that have to be monitored and that, if satisfied, cause a change from mode α to some new mode β .

A *normal mode* is one in which the system is functioning well in delivering its basic objectives. An *exceptional mode* is essentially an error recovery situation and is not considered a desirable operation. The aim is to solve the following:

Validation Problem. In designing an analogue-digital system with many modes we wish to validate – by testing and verification – that, relative to assumptions about the physical components,

1. There are sufficiently many modes to cover all cases, including exceptional modes; and
2. Given a subset of initial states, the behaviour of the system over a given time interval can be guaranteed to stay within the normal modes.

5 Designing a mode implementation

5.1 The architecture of a multimodal analogue-digital system

The architecture is designed to gather information about physical behaviour; monitor and evaluate the behaviour as a mode; and supervise a mode transition if needed. The components needed to make the software are summarised in the architecture depicted in Figure 4; compared with Figure 3, the physical system is omitted and, for simplicity, now only two modes α and β , and one transition from mode α to mode β , are shown. The next three sections will be devoted to explaining the principles behind this choice of components and architecture.

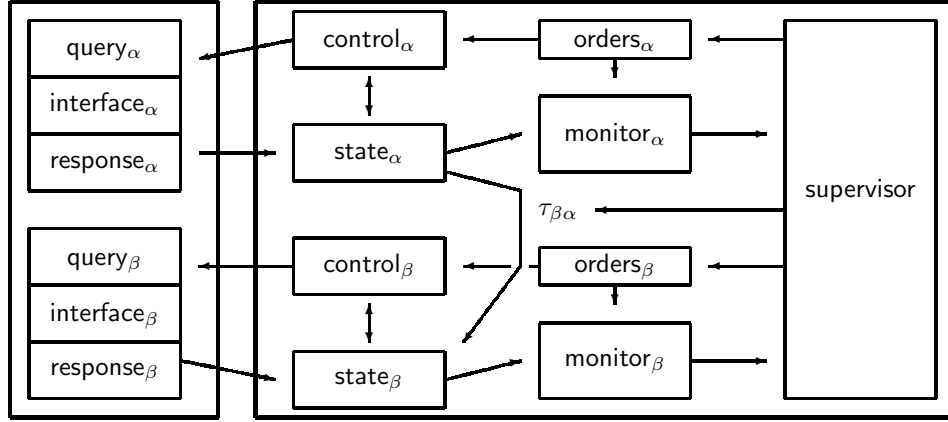


Figure 4: Details of a two mode analogue-digital system

5.2 The architecture of a mode

For each mode $\alpha \in \mathcal{M}$ we have the following components:

state_α: The data in this component constitutes mode α 's picture of the physical system, i.e., the state space State_α . It can be broadly divided into *history*, *observations*, *deductions* or *estimations*. Also system variables are likely to be accompanied by estimates of their errors, and the time of last measurement. Much of this data may be unused by **control_α**, but more of it is likely to be used by **monitor_α** (more on this later). **interface_α**: This is the only component of the program that communicates with the physical system. Its input to the interface from the rest of the program are queries from a set Queries_α , and lie in its querylist_α , being the list of tasks to be performed with the physical system. Queries are requests to make observations, and controlling instructions. The outputs from the interface to the rest of the program are the responses from a set Responses_α and are passed to the component **state_α**. Responses may be values it observes, and answers to questions. The interface has a default list of tasks designed to keep the states in **state_α** up to date. The tasks may be assigned various priorities, e.g., the tuples

- $\langle \text{single action, urgent priority, apply brakes} \rangle$
- $\langle \text{repeat hourly, measure temperature, accuracy } 1^\circ \rangle$
- $\langle \text{single action, standard priority, measure temperature, accuracy } 1^\circ, \text{ if previous measurement older than 1 minute} \rangle$

Here we see that some tasks may be more urgent than others, while some may be instructions to be repeated. They may come with conditions for the measurements to be made, or with additional information, such as a required accuracy.

Why is this rather indirect method used for communication? One answer is that the controlling algorithm is really not interested in the details of a particular measuring device. Should a sensor wear out and be replaced by one from a different manufacturer, only the interface would have to be updated.

control_α: This component of mode α actually makes decisions about controlling the physical system. It reads data from components **state_α** and **orders_α**, and writes to **state_α** (typically *estimations*) and the *query list* for **interface_α**. It might have some ability to respond to observed problems (e.g., by increasing the frequency of making some observations).

orders_α: While the response of **control_α** may be largely determined by the code, there may be some flexibility in its behaviour. This could be viewed as setting the overall strategy for mode α . Let **orders_α** be the set of possible instructions for mode α . A simple example would be temperature control for a building, where the inhabitants could set the desired temperature range, and this range would be put in as an instruction in **orders_α**.

Definition. A *physical mode* α is specified by a set $State_\alpha$ of states whose role is

- (i) *Monitoring*: to observe and act upon the behaviour of the system; and
- (ii) *Reflexivity*: to evaluate its own performance in faithfully representing that behaviour.

Various functions on $State_\alpha$ arise in specifying examples. For example, in monitoring, one expects functions to represent the generation of queries and the collection of responses:

$$Q : State_\alpha \times Order_\alpha \rightarrow Query_\alpha \text{ and } R : Response_\alpha \times State_\alpha \rightarrow State_\alpha.$$

In the Section 6 we will introduce the functions for reflexivity. Before that discuss reliability issues.

5.3 Accuracy and consistency

In an analogue system there are both measured and calculated data. Tracking errors due to measurement, approximation in calculation, time delays, or any other source is vital to the system. Errors are a measure of how well mode α is modelling reality, and they can tell if another mode of operation might do a better job in current circumstances. Changing modes is our primary problem, so how do we know that mode α is doing its job of coping with reality?

Imagine an analogue-digital system with a constant z of the system. Initially we measure $z \in (0, 5)$, and some time later we measure $z \in (4, 7)$. In what interval do we now know z to be? It is tempting to say $z \in (4, 5)$, but maybe ...

Wear. The sensor is degrading with time, and that an increase in the midpoint of the interval is part of a gradual drift?

Domain. The day of the second measurement was very hot, and the specification of the sensor says that it should be used at cooler temperatures; that caveat keeps the sensor engineers in the clear, but what about the measurement the system has to work with?

Calibration. The equipment was not recalibrated after a minor shock, and the constant z got shifted.

Now imagine these complicating factors are not known (e.g., because the system does not have the sensors to report them). Our deduction $z \in (4, 5)$ becomes

questionable. We may understand the errors, but do we understand the errors of the errors?¹⁰

Consider the worse case of a function $f(t)$ of time in an analogue-digital system. We measure $f(0) \in (0, 5)$, and the equations of motion given by a theoretical analysis of the system predict $f(1) \in (-2, 4)$. We then measure $f(1) \in (3, 6)$. We have all the previous complications (1)-(3) above, plus

Validity. Are the equations of motion really valid, or did we go outside their range of validity, or their numerical stability, or are there unexpected *unknown unknowns*?

Surely we want the measured value to be inside the range of the computed value? Measure $g(0) \in (0, 2)$, calculate $g(1) \in (-450, 900)$ and subsequently measure $g(1) \in (2, 4)$. The huge error margin in the calculated value simply illustrates that it is hopelessly inaccurate. The value of all calculated extrapolations tends to garbage as time increases, it is just a matter of how quickly.

It is important to distinguish between two sorts of *observed* errors, i.e., errors that we know about: If we estimate that the temperature should be in the range (23.1, 23.2) and it is observed to be in the range (23.4, 23.5) then we have an *inconsistency*. However, if we only need to know the temperature to an accuracy of one degree, then the estimate is *accurate*. In the example of $g(t)$ in the last paragraph, we had an example of *consistency* and *inaccuracy*.

Accuracy matters because being inaccurate may cause a failure of control for the system. Inconsistency matters because it tells us that there is a disagreement between the predicted behaviour and the observed behaviour of the system, and thus that the predicted behaviour should be considered unreliable.

6 Evaluating a mode implementation

6.1 The architecture of mode evaluation

To the four components introduced in subsection 5.2 for gathering information, for each mode $\alpha \in \mathcal{M}$ we now add a fifth for evaluation:

monitor _{α} : This measures the suitability of the model and is the only component of mode α which sends information to the system supervisor that decides on mode transitions.

Definition. The output of the component **monitor _{α}** is specified by a *mode evaluation function*

$$P_\alpha : \mathcal{M} \times \text{State}_\alpha \times \text{Orders}_\alpha \rightarrow [0, 1].$$

The number $P_\alpha(\beta, x, o) \in [0, 1]$ assesses how well mode $\beta \in \mathcal{M}$ *could* run the system, given the current picture $x \in \text{State}_\alpha$ of the system and the current

¹⁰ Compare the idealised situation in computable analysis. Take a computable real number y , and suppose we calculate $y \in (0, 5)$. Some time later we calculate $y \in (4, 7)$, so we deduce that $y \in (4, 5)$. As we perform more calculations, the error bound to within which we know y never increases.

orders $o \in \text{Orders}_\alpha$. The value of the function is normalised by two given values $0 < p_{\text{low}} < p_{\text{high}} < 1$ so that

$$\begin{aligned} P_\alpha(\beta, x, o) = 0 & \quad \text{means mode } \beta \text{ is } \textit{incompatible} \text{ with the current state;} \\ P_\alpha(\beta, x, o) < p_{\text{low}} & \quad \text{means mode } \beta \text{ models the system, but } \textit{inadequately}; \\ P_\alpha(\beta, x, o) > p_{\text{low}} & \quad \text{means mode } \beta \text{ models the system } \textit{adequately}; \\ P_\alpha(\beta, x, o) > p_{\text{high}} & \quad \text{means mode } \beta \text{ models the system } \textit{well}. \end{aligned}$$

Note that the function P_α is total. The component is responsible for assessing (i) observed problems, such as *consistency* and *accuracy* (Section 5.3), and (ii) the extent the observed state of the system is consistent with the objectives, including the current orders $o \in \text{Orders}_\alpha$.

6.2 Simplicial complexes and the geometry of the modes

Next we are concerned with the large scale behaviour of the system and, therefore, the relationships between the modes.

For some modes $\alpha, \beta, \gamma \in \mathcal{M}$, the overlap $\alpha \cap \beta \cap \gamma$ might be nonempty. This means, is it possible for the state of the physical system to be simultaneously in (say) all three modes α, β, γ . From such intersection information, we will build a geometric object, called the *nerve* of \mathcal{M} [3].¹¹ The nerve is a simplicial complex.

Simplicial complexes. A *0-simplex* is a point or vertex. A *1-simplex* is a line segment connecting two vertices. A simplicial complex is made of several simplices, so we can see that a simplicial complex consisting of 0-simplices and 1-simplices is just a graph. However, simplicial complexes generalise graphs by allowing higher dimensional constructions. A *2-simplex* is a triangle bounded by three 1-simplices. Figure 5 shows an example of a simplicial complex consisting of one 2-simplex, ten 1-simplices and eight 0-simplices; the 0-simplices are vertices v indexed by $\{\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \theta, \phi\}$.




Figure 5: A simplicial complex

The shading indicates that $\alpha\beta\gamma$ is a 2-simplex bounded by the 1-simplices $\alpha\beta$, $\alpha\gamma$ and $\beta\gamma$. Note that there is not a 2-simplex $\alpha\delta\gamma$. The dimensions continue to increase by a *3-simplex* being a tetrahedron (= triangular based pyramid) bounded by four 2-simplex sides, etc. Simplicial complexes are used in topology to give a “computable” presentation of topological spaces [35].

¹¹ For the general idea of a nerve of a category rather than a cover by subsets, see [34].

Coordinates for the simplicial complex are illustrated in Figure 6, which consists of one 2-simplex, four 1-simplices and four 0-simplices $\{\alpha, \beta, \gamma, \delta\}$ (the 2-simplex $\alpha\beta\gamma$ is not shaded for clarity).




Figure 6: Points in a simplicial complex

The 1-simplex $\gamma\delta$ consists of points

$$t_\gamma v_\gamma + t_\delta v_\delta \text{ for } t_\gamma, t_\delta \in [0, 1] \text{ and } t_\gamma + t_\delta = 1.$$

We identify the ends of the interval with the vertices, so $1 v_\gamma + 0 v_\delta$ is identified with v_γ and $0 v_\gamma + 1 v_\delta$ is identified with v_δ .

The 2-simplex $\alpha\beta\gamma$ consists of points

$$t_\alpha v_\alpha + t_\beta v_\beta + t_\gamma v_\gamma \text{ for } t_\alpha, t_\beta, t_\gamma \in [0, 1] \text{ and } t_\alpha + t_\beta + t_\gamma = 1.$$

We identify the sides of the 2-simplex triangle with the corresponding 1-simplex intervals, for example $\frac{1}{3}v_\alpha + \frac{2}{3}v_\beta + 0 v_\gamma$ is identified with $\frac{1}{3}v_\alpha + \frac{2}{3}v_\beta$ in the 1-simplex $\alpha\beta$.¹²

Note that we do not have $\frac{1}{4}v_\beta + \frac{1}{4}v_\gamma + \frac{1}{2}v_\delta$ in the simplicial complex in Figure 5, as there is no 2-simplex $\beta\gamma\delta$.

Construction of nerve. The nerve $N(\mathcal{M})$ of \mathcal{M} is a simplicial complex defined as follows.

For every mode $\alpha \in \mathcal{M}$, there is a point or vertex v_α . For two different modes $\alpha, \beta \in \mathcal{M}$ with $\alpha \cap \beta$ not empty there is a 1-simplex $\alpha\beta$.

For three different modes $\alpha, \beta, \gamma \in \mathcal{M}$ with $\alpha \cap \beta \cap \gamma$ not empty there is a 2-simplex $\alpha\beta\gamma$, and so on. To summarise this:

Definition. The element $\sum_{\alpha \in \mathcal{M}} t_\alpha v_\alpha$ is in the nerve $N(\mathcal{M})$ of set \mathcal{M} of modes if, and only if,

- (i) each $t_\alpha \geq 0$ and $\sum_{\alpha \in \mathcal{M}} t_\alpha = 1$, and
- (ii) taking the subset of $\alpha \in \mathcal{M}$ for which $t_\alpha > 0$, the intersection of all the modes in the subset is nonempty (i.e., the subset gives a simplex of appropriate size).

¹² This sort of coordinate system for a triangle may be compared with the RGB colour system.

6.3 The mode evaluation and classification functions

How can there be uncertainty as to which mode the system is in? Near the boundaries between two subsets it would be common for there to be an overlap. Suppose the program is in mode $\alpha \in \mathcal{M}$. Let mode α have mode evaluation function

$$P_\alpha : \mathcal{M} \times \text{State}_\alpha \times \text{Orders}_\alpha \rightarrow [0, 1].$$

Recall from subection 6.1, that $P_\alpha(\beta, x, o)$ is a measure of how well $x \in \text{State}_\alpha$ could be described by mode $\beta \in \mathcal{M}$ given the orders $o \in \text{Orders}_\alpha$; and for value $p_{\text{low}} > 0$, if $P_\alpha(\beta, x, o) > p_{\text{low}}$ then $x \in \text{State}_\alpha$ could be reasonably described by mode β , given orders $o \in \text{Orders}_\alpha$.

Definition. The *mode classification function*

$$N_\alpha : \text{State}_\alpha \times \text{Orders}_\alpha \rightarrow N(\mathcal{M})$$

is defined as follows:

$$N_\alpha(x, o) = \frac{\sum_{\beta: P_\alpha(\beta, x, o) > p_{\text{low}}} (P_\alpha(\beta, x, o) - p_{\text{low}}) v_\beta}{\sum_{\beta: P_\alpha(\beta, x, o) > p_{\text{low}}} (P_\alpha(\beta, x, o) - p_{\text{low}})}. \quad (1)$$

The mode classification function *sends the current state x and orders o to a point in the nerve, to give a geometric picture of the situation with respect to the modes.*

Note that we only sum over modes β for which $P_\alpha(\beta, x, o) > p_{\text{low}}$, i.e., we only consider those modes which have been certified as being reasonable descriptions of the system, given its observed state. Of course, an objective of the design of the system is that for *any* possible state of the system there is a mode that describes it adequately, and so that there is always a consistent choice of such modes.




Figure 7: Mode classification for a four mode analogue-digital system

As Figure 7 illustrates, we are considering systems for which there is no unified or global model; the nerve binds the modes together.

6.4 Exceptions: known unknowns and unknown unknowns

What could possibly go wrong? Many things could go wrong. Recalling *Principle 0: Robustness*, algorithms ought to respond to *any* eventuality. It is useful to consider Rumsfeld's classification of unknowns. In our situation, and perhaps in system design generally, there are natural working definitions:

Definition. A *known unknown* is an event that we knew might happen, we knew the approximate form of the event, and, specifically, that there is software to handle the event. An *unknown unknown* is any event that we did not expect and we do *not* have software to handle.

For an example of a known unknown, in the robotic car race of subsection 3.1, another car could unexpectedly join the race after it has started; a designer can include a provision for adding cars during the race, which just requires some details to be added if and when the event occurs.

How do we know when an *unknown unknown* is happening? The two mathematical singularities that can occur with the mode classification formula (1) give a good indication of an observed *unknown unknown*:

Partiality. The set of modes β with $P_\alpha(\beta, x, o) > p_{\text{low}}$ may be empty, i.e., there are no modes which, in the opinion of the designers, model the current state adequately. Since we assumed $p_{\text{low}} \neq 0$, we have some room in which we may choose the best of a lot of bad choices. This is discussed further in Subsection 7.4.

Contradiction. For the modes in Figure 5, a classification value of $N_\delta(x, o)$ of $\frac{1}{2}v_\gamma + \frac{1}{2}v_\delta$ is fine, as there is a line connecting the vertices v_γ and v_δ , so the designers knew that those modes could in principle be consistent at a given time. However, the value $\frac{1}{2}v_\beta + \frac{1}{2}v_\delta$ is invalid, it does not lie in the simplicial complex. This geometric property indicates that the designers never considered the possibility that these modes v_β and v_δ could both be valid descriptions of the same physical state. We have a possible contradiction – this is discussed further in Subsection 7.5.

7 Making transitions between modes

7.1 The architecture of mode transition

There are two components that accomplish a change of mode, one of which was not illustrated earlier in Figure 4:

supervisor: This component decides on the change of mode from (say) mode α to mode β . Its input is the value of the mode evaluation function $P_\alpha(\beta, x, o)$ from **monitor** $_\alpha$. It assigns the mode β to control the system by

- (i) changing state information State_α , and
- (ii) modifying that mode's Orders_α .

In the case of observed unknown unknowns it must decide on the system's response (e.g., partiality and consistency – see subsections 7.5 and 7.4).

The component is the only part of the program which receives commands from outside agencies; for example, it is responsible for initialising the whole system on startup; and it can receive an over-ride to move to an emergency shutdown mode.

transfer_α: This component moves information between modes – changing states – and its specification is the family of *mode transition functions*:

$$\tau_{\beta\alpha} : \text{State}_\alpha \times \text{Orders}_\alpha \rightarrow \text{State}_\beta \times \text{Orders}_\beta$$

It is the only part of mode α which needs to know anything about State_β . The mode transition function is partial. When the function $\tau_{\beta\alpha}$ is called by the supervisor and computed, control is transferred to mode β .

7.2 The mode transition functions $\tau_{\beta\alpha}$: commenatry

If we decide to move from mode α to mode β , how do we actually implement this? Each mode α carries a picture state_α of the physical system, so to move from α to β we need a mode transition function $\tau_{\beta\alpha}$, so that our picture of reality is redrawn in a form understood by mode β . However, this function need not be defined on all of State_α ; in fact the designers are only likely to have implemented the function $\tau_{\beta\alpha}(x)$ on the condition that β is an adequate or good description for $x \in \text{State}_\alpha$.

The transition functions also depend on the standing orders Orders_α . For example, when controlling temperature, we might simply store the target temperature of a building in the standing orders, e.g., $\text{Orders}_\alpha = \{21^\circ\text{C}\}$. In calculating the mode evaluation function P_α we obviously need to take account of the target temperature, but why is it needed in $\tau_{\beta\alpha}$? Simply copying the standing orders, so we would get $\text{Orders}_\beta = \{21^\circ\text{C}\}$ may negate the point of changing: the target temperature could be modified by various factors, such as the time of day or the cost of energy, giving e.g., $\text{Orders}_\beta = \{19^\circ\text{C}\}$.

Thus, we need the mode transition function to be total when restricted to the form

$$\tau_{\beta\alpha} : \{(x, o) \in \text{State}_\alpha \times \text{Orders}_\alpha : P_\alpha(\beta, x, o) > p_{\text{low}}\} \rightarrow \text{State}_\beta \times \text{Orders}_\beta. \quad (2)$$

We need to assume that the designers have done a sufficiently good job of implementing $\tau_{\beta\alpha}$ so that the system can carry on from this position in mode β . If this proves too difficult to do, then it is possible that β is not a reasonable description for all $(x, o) \in \text{State}_\alpha \times \text{Orders}_\alpha$ with $P_\alpha(\beta, x, o) > p_{\text{low}}$, and that the function P_α may have to be redefined.

Note that not all the state variables in State_β need be set by the mode transition function in (2), some may be left undefined. To use the example of

the car race in subsection 3.1, if we are in mode α we may decide to switch to mode β because another car is within a certain distance. However there is no reason why we needed to have much information on that particular car earlier. In particular, we may have no idea of its velocity, or its identity, so these may be left as undefined after the switch in modes. In this case, it is to be hoped that the designers of mode β have given a high priority to finding a value for these undefined variables.

One problem with changing modes is what happens when it is done several times in quick succession. Suppose that we have just used the function in (2) to move from mode α to β . It may happen that we almost immediately switch to mode γ . However, the designers could only be expected to have taken account of this possibility if γ is a reasonable description of $\tau_{\beta\alpha}(x) \in \text{State}_\beta$. In that case, we get the composition

$$\tau_{\gamma\beta} \circ \tau_{\beta\alpha} : \{y \in \text{State}_\alpha \times \text{Orders}_\alpha : P_\alpha(\beta, y) > p_{\text{low}} \text{ and } P_\beta(\gamma, \tau_{\beta\alpha}(y)) > p_{\text{low}}\} \rightarrow \text{State}_\gamma \times \text{Orders}_\gamma. \quad (3)$$

However, we might have moved from α to γ in one go, using

$$\tau_{\gamma\alpha} : \{y \in \text{State}_\alpha \times \text{Orders}_\alpha : P_\alpha(\gamma, y) > p_{\text{low}}\} \rightarrow \text{State}_\gamma \times \text{Orders}_\gamma. \quad (4)$$

As we could have taken either path from α to γ , we should hope that there is a reasonable consistency between $\tau_{\gamma\beta} \circ \tau_{\beta\alpha}$ and $\tau_{\gamma\alpha}$ as functions on the subset of $\text{State}_\alpha \times \text{Orders}_\alpha$ where all three inequalities in (3) and (4) are satisfied; this should be part of the design requirements.

7.3 Normal running

Changing modes changes the way a system is viewed and may be needed to simplify and frame a computational description. Changing the way a system is viewed comes with a cost: change takes time, it may introduce inaccuracy or inconsistency, and it may lose information. For a robust many mode system, a method is needed by which the supervisor minimises the changes of mode. Periodically, the program checks that it is in a correct mode to describe the system, as follows:

If we are in mode α , and $P_\alpha(\alpha, x, o) > p_{\text{high}}$, then $x \in \text{State}_\alpha$ together with the standing orders $o \in \text{Orders}_\alpha$ should be well described by mode α . We therefore remain in mode α .

If $P_\alpha(\alpha, x) \leq p_{\text{high}}$ we calculate the other values of $P_\alpha(\beta, x, o)$ and use formula (1) to find $N_\alpha(x, o)$. We could then select the $\beta \in \mathcal{M}$ which gives

$$P_\alpha(\beta, x, o) = \max_{\gamma} \{P_\alpha(\gamma, x, o)\}.$$

We could modify this in the event of several large values by using the more detailed geometry of the symplectic space by choosing a vertex v_β which $N_\alpha(x, o)$ is approaching, comparing the last several values of $N_\alpha(x, o)$. In this manner we

pick a mode which the state is entering, and so might be expected to stay in for longer.

Recalling the *Principle 0: Robustness*, we attempt to implement a *robust* system, with the following working definition:

Definition. The system of modes is *robust* if there are a sufficient number of modes so that we can always choose a mode which models the state well, i.e., with mode evaluation function $> p_{\text{high}}$.

The system of modes is *adequate* if there are a sufficient number of modes so that we can always choose a mode which models the state adequately, i.e., with mode evaluation function $> p_{\text{low}}$.

These notions reflect the idea that the modularisation of physical behaviour covers all eventualities, normal and exceptional. We can attempt to *verify* that this is the case, by axiomatising physical assumptions about the behaviour of the system and the reliability of our sensors and actuators. However, as von Moltke observed, ‘no plan survives contact with the enemy’: our enemy is the physical reality of the system. We need to know when the plan is going wrong, and what to do. In summary, *the plan may fail but the real world continues to move on and decisions have to be taken*.

We examine the two exceptions which can be thrown by the mode transition function N_α , introduced in subsection 6.3.

7.4 Partiality exception: Into the unknown

In the case of the *partiality exception* in subsection 6.3 we have the problem:

$$\{\beta \in \mathcal{M} : P_\alpha(\beta, x, o) > p_{\text{low}}\} = \emptyset.$$

In other words, for all β we have $P_\alpha(\beta, x, o) \leq p_{\text{low}}$, i.e., the software does not think that any mode β reasonably describes the picture and orders (x, o) .

The assumption, in Section 6.2, that $p_{\text{low}} > 0$ means we have some ‘wiggle room’, a gap between where ‘reasonable description’ ends and ‘complete garbage’ begins. We can still try to choose the best of a lot of bad choices – we may take the largest value of $P_\alpha(\beta, x, o)$ even if it is $\leq p_{\text{low}}$. Essentially we may have a mode β which the engineers thought was some sort of description of the system, but they were not prepared to certify as a ‘reasonable description’. When using such a mode β to control the system, it may be wise to have a *safety strategy* for β – a version of Orders_β written on the understanding that mode β , while the best available, may still not be a reliable representation of the real world.

But which β do we choose? It is possible that the values of $P_\alpha(\beta, x, o)$ may not distinguish well between values in the range $\leq p_{\text{low}}$ – after all, the system was never supposed to be run in that range. When lost but with a last known position mode α , we can bias the search for new modes towards “nearby” modes. The simplicial complex provides an idea of the distance between modes: the shortest path in terms of number of edges or 1-simplices between the vertices.

For example, if for $\alpha, \beta \in \mathcal{M}$ we have a γ with $\alpha \cap \gamma \neq \emptyset$ and $\gamma \cap \beta \neq \emptyset$ but $\alpha \cap \beta = \emptyset$ then $d(\alpha, \beta) = 2$. (Alternatively, if the collection of modes \mathcal{M} is part of a multi-scale hierarchy, we might move to a coarser collection of modes: see subsection 11.3.)

If the system does not swiftly return to a known mode, and the system has not collapsed, a warning would alert people or another system to intervene.¹³

7.5 Contradiction exception

In the case of the *contradiction exception* in subsection 6.3 we have the problem:

$$\{\beta \in \mathcal{M} : P_\alpha(\beta, x, o) > p_{\text{low}}\} \neq \emptyset \text{ but has no common intersection.}$$

In the case of three modes β, γ, δ with $P_\alpha(\beta, x, o) > p_{\text{low}}$, $P_\alpha(\gamma, x, o) > p_{\text{low}}$ and $P_\alpha(\delta, x, o) > p_{\text{low}}$, the following two configurations of intersections of modes in Figure 8 both give problems:




Figure 8: Intersections of modes

We have statements that $(x, o) \in \text{State}_\alpha \times \text{Orders}_\alpha$ could be reasonably described by a collection of modes which has no common intersection. But surely this is a real contradiction, as (x, o) would be in the intersection? The answer is no, and the reason is lack of information or errors. Remember that the point of modes is that state_α is *not* the state space of the whole system, but rather a picture created by the software to describe one aspect, mode α , of the system. Let us digress with a new scenario to illustrate this:

Example: Aircraft. A military aircraft spots a new contact on its radar screen. There is, as yet, insufficient information to say if it is a hostile military plane or a civilian plane. Both possibilities *could* reasonably describe the situation as either might well be true or false; but, in particular, there is no possibility that describes the situation well. If we just have the two disjoint modes $\alpha = \text{hostile}$ and $\beta = \text{civilian}$ to choose from, we have the case where there is more than one possible mode assignment consistent with the data (not necessarily a problem), but the assignments are, in reality, mutually exclusive (a problem). The reason is lack of data. We know that the designers did not consider this case, as if so they would have drawn a line joining α and β in the simplicial complex, and no

¹³ The required abstract understanding may well be beyond current artificial intelligence; see [20] for progress on this front.

exception would have occurred. Of course, it might be said that the designers *should* have considered this case, and that would be a valid point to make at a committee of enquiry into any disaster resulting from the situation, but is not much help at the time.

So we know there is a problem, but what do we do? We can flag an error, and hope that a human puts the matter right before something unfortunate happens – if the system can simply be shut down in a safe mode. Extra information could resolve the problem, and the exception should trigger several queries for more information. But what do we do *now*? Here are four options that might work in the airplane scenario:

Hawk. Knowing that there is a significant chance it might be wrong, the software takes the most likely mode, i.e., the largest P_α , and, e.g., as a result fires missiles to shoot the target down. For example, in the absence of other information, the matter was decided because the unknown plane was not on a commercial flight path.

Dove. Again knowing that there is a significant chance it might be wrong, the software takes the most likely mode. But although the plane may assume that the other aircraft is hostile α , it is operating under safety strategy where standing orders are used when the mode is active but considered to be unreliable. In this case an obvious thing to write in Orders_α would be ‘do not fire missiles’.

Fail-safe. Written into the code for P_α is the guarantee:

$$P_\alpha(\text{civilian}, x, o) > p_{\text{low}} \implies P_\alpha(\text{hostile}, x, o) < P_\alpha(\text{civilian}, x, o).$$

This means that a taking the largest choice defaults to civilian rather than hostile, in the case where the situation is reasonably described by the unknown plane being civilian. Thus, if we fail to identify the plane successfully, the default is fail-safe (well, as far as the possible civilians are concerned). Of course there is a slight inconsistency here – if the engineers rigged this fail-safe, why did they not simply include code for handling an intersection of hostile and civilian?

Fail-safe: Consensus. During the design process the two design teams for the modes are asked: ‘Suppose that we have an inconsistent mode allocation, but the evidence indicates that your mode may be applicable. What orders should be issued as a result, given that your mode may equally well not be applicable?’. The civilian mode designers might reply ‘do not fire missiles’, and the hostile mode designers might reply ‘take evasive action to avoid attack’. Now whichever mode takes control of the system is given *consensus* orders containing both the orders above. This is effectively a way of making up some sort of fail-safe strategy as the design develops.

8 Systems modelled by manifolds and geodesics

Many physical systems are modelled by dynamical systems based on differential manifolds [30,4,5,36]. A manifold is a topological space whose neighbourhoods have local coordinate systems based on \mathbb{R}^n . These local coordinate systems are

called *charts*. In fact, a manifold is *determined* by its charts. Each chart is a coordinate system optimised for the description of a dynamical behaviour in its neighbourhood. The behaviour of the dynamical system in time is a path through the manifold passing through the charts. Thus, the manifold is a mathematical construction that combines distinct local representations into a global representation of states and behaviours.

If a complex system can be modelled by a manifold then we are in a special case where there are many tools that can be applied to analyse the system. We will show how

- (i) modes can be based upon the local representations by charts, and
- (ii) the modes change as the behaviour of the dynamical system changes in time.

Such dynamical systems are common. They are, however, an important example as our theory of multi-modal analogue digital systems is specifically designed so as *not* require there to be a global state space for the physical system.

8.1 Smooth manifolds

A topological space X is an n -dimensional smooth manifold if it possess a family of charts. Let \mathcal{M} be some index set. According to the usual definition, a *chart* consists of

$$\text{open subsets } U_\alpha \subset X \text{ and homeomorphisms } \phi_\alpha : U_\alpha \rightarrow V_\alpha,$$

where each V_α is an open subset of \mathbb{R}^n and $\alpha \in \mathcal{M}$.

Furthermore, for every $\alpha, \beta \in \mathcal{M}$ with $U_\alpha \cap U_\beta$ not empty, the *transition function*

$$\tau_{\beta\alpha} = \phi_\beta \phi_\alpha^{-1} : (\phi_\alpha(U_\alpha \cap U_\beta) \subset V_\alpha) \rightarrow (\phi_\beta(U_\alpha \cap U_\beta) \subset V_\beta) \quad (5)$$

is a smooth function (i.e., differentiable arbitrarily many times). Further if $U_\alpha \cap U_\beta \cap U_\gamma$ is not empty we have the equality

$$\tau_{\gamma\alpha} = \tau_{\gamma\beta} \tau_{\beta\alpha} : (\phi_\alpha(U_\alpha \cap U_\beta \cap U_\gamma) \subset V_\alpha) \rightarrow (\phi_\gamma(U_\alpha \cap U_\beta \cap U_\gamma) \subset V_\gamma) . \quad (6)$$

To compute with this we need a function to tell us which chart we are in. This is given by a *partition of unity* [32] p. 40, which is a collection of functions:

$$\theta_\alpha : X \rightarrow [0, 1] \text{ with } \sum_\alpha \theta_\alpha(x) = 1 \text{ for all } x \in X \text{ and } \theta_\alpha(x) = 0 \text{ if } x \notin U_\alpha.$$

(Thus we have a ‘part’ of ‘unity’ 1 for each open set.) We relate these functions to the charts $V_\alpha \subset \mathbb{R}^n$ by defining

$$P_\alpha(\beta) = \theta_\beta \phi_\alpha^{-1} : V_\alpha \rightarrow [0, \infty) . \quad (7)$$

Figure 9 illustrates the nerve for a covering of a manifold with charts labelled by $\{\alpha, \beta, \gamma, \delta\}$ (the triangle is not shaded for clarity). Keeping this order, the function $(P_\alpha(\alpha), P_\alpha(\beta), P_\alpha(\gamma), P_\alpha(\delta))$ gives a function into the nerve, for illustration

mapping the displayed point on the manifold (in α coordinates) to the point on the nerve.

For purposes of computation it is convenient to extend this $P_\alpha(\beta) = \theta_\beta \phi_\alpha^{-1} : V_\alpha \rightarrow [0, \infty)$ to a function on domain \mathbb{R}^n by taking the function to be zero outside V_α .




Figure 9: The nerve for a manifold

If there are originally m open sets, then if we choose $p_{\text{low}} = \frac{1}{m+1}$, for any α and $y \in V_\alpha$ we are guaranteed to have at least one β so that $P_\alpha(\beta, y) > p_{\text{low}}$.¹⁴ Recalling the definition in subsection 7.3, we have:

Lemma. *The construction of an analogue-digital system with finitely many modes based upon the finitely many charts of a differential manifold as specified above is adequate.*

The theory of manifolds is extensive and contains many intriguing results about inequivalent families of charts and the embedding of n -dimensional manifolds in sufficiently higher dimension spaces \mathbb{R}^N . The question as to whether and how computation on manifolds should be split into coordinate charts – rather than, e.g., regarding it as a subset of a higher dimensional Euclidean space – is answered in the extensive and closely related literature on the choice of grids for global climate models, e.g., [2].

8.2 A dynamical system

Consider how a practical problem involves the idea of a manifold. The manifold represents numerically and geometrically the space of all possible states of a system. The behaviour of the system in time is represented – geometrically and numerically – by a vector field and paths in the manifold. There is a differential equation governing the time evolution of the coordinates of states, which are points in the path traversing the manifold. Periodically, we have to *change* coordinate charts.

For example, consider the idea of a *geodesic* or minimum distance path on a curved surface; it has been likened to the path of an ant crawling on the

¹⁴ A better bound for p_{low} in terms of the dimension is possible from the Čech-Lebesgue covering dimension [19].

surface of an apple [27]. In Einstein's general theory of relativity [16] the paths of the planets are given by geodesics in a space-time manifold, and the resulting prediction of the perihelion precession of Mercury was central in convincing people about that theory. (See [27,31] for the mathematical details.)

The definition of a manifold and geodesic motion fits our idea of computation of many modes as follows. To begin, we assume that we have a register machine with real number registers, and operations that enable it

(i) to calculate precisely all those inconvenient transition functions between charts; and even

(ii) to solve the differential equations for the geodesics.

The modes are exactly the set of charts indexed by \mathcal{M} . The picture of the system corresponding to each mode $\alpha \in \mathcal{M}$ is $\text{State}_\alpha = \mathbb{R}^n \times \mathbb{R}^n$. We consider $(y(t), \dot{y}(t)) \in \mathbb{R}^n \times \mathbb{R}^n$, where $y(t)$ is the state (e.g., position of our ant, airplane or planet) at time t in the coordinate chart V_α , and $\dot{y}(t)$ is its rate of change or velocity. No attempt to place any restriction forcing $y(t) \in V_\alpha$ is needed: it is the job of the functions $P_\alpha(\beta, (q, p)) = \theta_\beta \phi_\alpha^{-1}(q)$ to ensure that charts are changed as appropriate. (This type of example does not include Orders_α as the program only has one purpose, with no variability.)

Next, we consider how these functions are computed. In computable analysis, computable reals and computable functions of computable reals are computed by programs generating rational numbers with error bars [41]. To find the position at a given time to a given error, we have to find how accurately we need to know the initial point, and then time step into the future, keeping track of errors and (rather messily) changing charts where we have to. There may be singularities, places beyond which we cannot extend the geodesics, but bearing that in mind, this is a perfectly consistent computable mathematical theory. However, we no longer have $\text{State}_\alpha = \mathbb{R}^n \times \mathbb{R}^n$, we need a more complicated data type incorporating error bounds for $(y(t), \dot{y}(t))$. (For analogue-digital systems see [18].)

What complications can arise between our mathematical model and the observations?

Overdue computation. The real physical system waits for no computer, it evolves regardless of whether the computer has finished a particular calculation. In computable analysis, if the velocity becomes too large we can reduce the time-step in the program to compensate. In real physical time we may not have this option, and the velocity may become so large that we even lose track of which chart or mode the system is in.

Chaotic computation. Many physical systems are chaotic, effectively meaning that even if we had an arbitrarily large amount of time to compute and an exact physical theory, we still could not measure the initial conditions exactly enough to extrapolate beyond a certain time limit. Of course, this could be a *known unknown*, and we can force the system to compensate by regularly observing the most chaotic variables, and using these to correct the calculated values. In the case of the solar system, this would be the positions of the planets around their

orbits, as opposed to the much more slowly changing orbital parameters (for the solar system see [21]).

9 A racing track

9.1 A single car

Consider a toy car race around the track shown in Figure 10:




Figure 10: A racing track

Mode α : Free running: For a single car, the slots on the track keeps the car moving in a given path, and all we have to worry about is getting the fastest time round the track, taking account of the problem that the maximum speed of the car is too fast for the semicircular parts of the track, so the car has to slow down there. For car $i = 1, 2$ this mode α^i has $\text{State}_{\alpha}^i = (q_i, v_i)$, where q_i is the position of car i and v_i is the velocity of car i .

Mode β : Near chicane: When two cars 1 and 2 are on the track, they can collide where the track narrows at the chicane, at position q_c . To stop this happening, a new mode β^i is introduced for car i for collision avoidance. We need data on both cars, so $\text{State}_{\beta}^i = (q_1, v_1, q_2, v_2)$.

Mode evaluation: For simplicity we assume that mode β^i is activated any time that car i is near the chicane. We have *mode evaluation functions*

$$\begin{aligned} P_{\alpha}^i(\alpha^i, (q_i, v_i)) &= f(|q_i - q_c|) , \\ P_{\alpha}^i(\beta^i, (q_i, v_i)) &= 1 - f(|q_i - q_c|) , \\ P_{\beta}^i(\alpha^i, (q_1, v_1, q_2, v_2)) &= f(|q_i - q_c|) , \\ P_{\beta}^i(\beta^i, (q_1, v_1, q_2, v_2)) &= 1 - f(|q_i - q_c|) , \end{aligned}$$

where $|q_i - q_c|$ is the distance from car i to the chicane, d is chosen according to the stopping distance of the car, and f is given by the graph in Figure 11:




Figure 11: The graph of f

The nerve: This consists of the 0-simplices α and β , and the 1-simplex (α, β) .




Figure 12: Single car nerve for the car race

Transition functions: We suppose that Orders is trivial for simplicity. Then

$$\tau_{\alpha,\beta}^1 : \text{State}_{\beta}^1 \rightarrow \text{State}_{\alpha}^1 \text{ is given by } \tau_{\alpha,\beta}^1(q_1, v_1, q_2, v_2) = (q_1, v_1) \text{ and}$$

$$\tau_{\beta,\alpha}^1 : \text{State}_{\alpha}^1 \rightarrow \text{State}_{\beta}^1 \text{ is given by } \tau_{\beta,\alpha}^1(q_1, v_1) = (q_1, v_1, ?, ?).$$

Here ? denotes an undefined value, and one which it is a high priority to find out about.

9.2 Several cars

Implementing the safety measures in the car race depends on knowing the positions of the other cars. This is likely to involve other processors, as competing teams would likely be unhappy at the thought of a single process controlling both team's cars. In other words, we have a network. Just how much cooperation between nodes of the network do we need to ensure that our safety measures work?

Autonomous: A team fitted their car with a camera to observe nearby cars. They can now implement mode β entirely independently of what any other team does.

Communal: The race organisers have made a condition of entry to the competition, that all cars must make their position public if they are near the chicane, in order to avoid accidents. If car 1 enters mode β^1 it queries car 2, and if car 2 is also in mode β^2 it is obliged to respond.

Priority: A particular team has incorporated into mode β software to maximise their team points in the competition, in the case that two of its cars are racing each other. For example, one car could let the other always go first into the chicane. If possible, the team might actually combine the processes β^1 and β^2 into a single process β^{12} running on a single computer.

10 A solar system

10.1 Real time modelling of a solar system

The Sun, Moon and Earth form a three body gravitational system, and there is no exact general solution known to the three body gravitational problem. The point is that we do not have a general three body problem with the Sun, Moon and Earth: To a very good approximation the Moon orbits the Earth, and the Earth-Moon system orbits the Sun, and our computational problem is vastly simplified.

More generally, take an example of a solar system with three planets $\{1, 2, 3\}$ orbiting a star. As long as the planets are a reasonable distance apart, the dominant feature in their orbits is the star's gravity, with minor influences from the other planets. This is modelled by processes or modes $\{\alpha^1, \alpha^2, \alpha^3\}$, which can easily be run on different processors or nodes, as occasional communication of the other planets orbits is sufficient to keep the planet i 's orbit updated.

Modes: Let us concentrate on planet 1. If one of the other planets, say 2, comes within a certain distance of 1, then we require much more detailed knowledge of the position and velocity of planet 2 to accurately model the path of planet 1. We will call this mode β_2^1 (reserving the superscript i for a process run on node i). Likewise there is a mode β_3^1 used when planet 3 is close to planet 1, and a mode γ_{23}^1 when both other planets are close to 1. Thus node 1 has four modes of operation, $\mathcal{M}^1 = \{\alpha^1, \beta_2^1, \beta_3^1, \gamma_{23}^1\}$.

Mode evaluation: To choose between these modes, we use indicator functions $\{w_2^1, w_3^1\}$. The definition of these is quite simple, the value $w_i^1 \in [0, 1]$ and $w_i^1 \cong 0$ when planet i is 'far' from planet 1, and $w_i^1 \cong 1$ when planet i is 'close' to planet 1. Thus w_i^1 may be thought of as a fuzzy logic truth value of the statement *planet i is close to planet 1*.




Figure 13: Configurations of planets

The indicator functions (w_2^1, w_3^1) are converted to mode evaluation functions by

$$(w_2^1, w_3^1) \mapsto ((1 - w_2^1)(1 - w_3^1), w_2^1(1 - w_3^1), (1 - w_2^1)w_3^1, w_2^1w_3^1), \quad (8)$$

where we use the order $\{\alpha^1, \beta_2^1, \beta_3^1, \gamma_{23}^1\}$. In Figure 13 we have a configuration of planets with, say as indicated by the dot on the square of (w_2^1, w_3^1) values, $(w_2^1, w_3^1) = (0.8, 0.2)$, giving mode evaluation functions $(0.16, 0.64, 0.04, 0.16)$. Thus the most suitable mode for node 1 in Figure 13 is β_2^1 . Figure 13 also shows the value of the mode evaluation functions for the corners of the square of (w_2^1, w_3^1) values.

The nerve: This is the solid tetrahedron with vertices $\{\alpha^1, \beta_2^1, \beta_3^1, \gamma_{23}^1\}$. It is a 3-simplex because, as shown in the example, we expect that all four mode evaluation functions can be nonzero simultaneously.

Transition functions: We have not specified the data structures associated to the modes, however if the modes β used Einstein's general relativity and mode α used Newton's theory, it might be imagined that quite a bit of data conversion would have to be done.

10.2 Cooperation between processes tracking planets

It was found that communication problems between different machines meant that running two separate processes for planets which were close together was too inaccurate (evidently this star system is more interesting and dangerous than ours). It was decided to instigate a cooperative process $\mathcal{M}^{12} \subset \mathcal{M}^1 \times \mathcal{M}^2$, consisting of just one mode $\beta^{12} = (\beta_2^1, \beta_1^2)$ running on one machine. Thus when node 1 and node 2 agree that planets 1 and 2 are close, but that planet 3 is further away, an attempt will be made to start a single process β^{12} to model the positions of planets 1 and 2. Should that attempt fail (for example a necessary communications link not be established), then the separate processes β_2^1 and β_1^2 will continue to run.

The cooperative process β^{12} ends when we leave $\mathcal{M}^{12} \subset \mathcal{M}^1 \times \mathcal{M}^2$. The process β^{12} makes sure that the pictures $\text{State}_{\beta_2^1}^1$ and $\text{State}_{\beta_1^2}^2$ in nodes 1 and 2 respectively are updated, and then switches control back to the individual nodes in modes $(\beta_2^1, \beta_1^2) \in \mathcal{M}^1 \times \mathcal{M}^2$. Given that (β_2^1, β_1^2) is no longer the preferred mode, either node 1 will then instigate a mode transition, or node 2 will do so. The cooperative mode transfers control to $(\beta_2^1, \beta_1^2) \in \mathcal{M}^1 \times \mathcal{M}^2$ simply because it was not necessary to duplicate the machinery of the mode transition functions.

Of course, we should also have $\mathcal{M}^{13} \subset \mathcal{M}^1 \times \mathcal{M}^3$ and $\mathcal{M}^{23} \subset \mathcal{M}^2 \times \mathcal{M}^3$ operating on the same principle. But what happens if all three planets are close? Following from two node cooperations, we could have three node cooperations $\mathcal{M}^{123} \subset \mathcal{M}^1 \times \mathcal{M}^2 \times \mathcal{M}^3$, consisting of just one mode $\gamma^{123} = (\gamma_{23}^1, \gamma_{13}^2, \gamma_{12}^3)$.

11 Concluding remarks and further directions

We have introduced the idea of a general analogue-digital system with many physical modes and created a data type, the nerve, to manage the relationship between the modes. We have discussed how modes arise and cover – or fail to cover – every physical eventuality. Finally, we return to the comparison of our approach with some established formal methods, and comment on the some

further refinements to do with classic questions about parallelism, hierarchy and security.

11.1 Comparison with hybrid systems

Our fresh look at analogue-digital systems introduces some new ideas for theorising hybrid systems. We believe our approach is distinct from the established computational theory of hybrid systems in a number of ways.

AD-interface. We think about systems that combine the analogue with the digital in very broad terms; we are not focussed on applications involving embedded systems, for example. We model from first principles, i.e., developing properties and an agenda from fresh reflections.

Oracles. We assume the digital component is computationally universal and can apply the full power of Church-Turing Thesis. This enables us to use computability theory and complexity theory in developing theory. Indeed, a central notion is that of a *physical oracle* to an algorithm, after Turing.

Measurement. We assume the analogue component generates physical measurements for the digital component to process. Ideally, measurements are continuous and are generated continuously, but the data processed are approximate and usually sampled discretely. Thus, the analogue component delivers streams of approximate data, which involve errors. We explicitly discuss data that may not match the mathematical models used to design software, or arise without any model available.

Modes. The systems are sufficiently complicated to require modularisation into distinct modes of physical behaviour. The distinct algorithms therefore need to change modes. Because of errors of measurement, and other factors, this mode switching is a complex nondeterministic phenomenon. We model the nondeterminism by paths through finite dimensional geometric objects, which we call nerves. A control sequence is a path through the modes of the system. The variations in the path are determined by the nerve. In the hypothetical global state space of the system, as physical conditions vary we consciously modify the control path. The nerve is a finite combinatorial representation of this, the higher dimensional structure of the nerve determines the transitions between control sequences.

11.2 Networks and cooperation

The two examples of analogue-digital systems with many modes in Sections 9 and 10 illustrate cooperation between analogue-digital systems: there are two cars and three bodies. In both cases the component systems are identical, but this is not necessary. In a quarry we could have diggers with modes {digging, moving} and trucks with modes {loading, unloading, moving}.

Consider a general network of analogue-digital systems and suppose that some of these systems can cooperate; the systems are nodes in the network. There are penalties for cooperation, such as communication speed limits between nodes of the network, or problems with the increased complexity of the system. Designers need to decide under what circumstances to attempt cooperation, and how to proceed if an attempt fails. There has to be a good reason to cooperate between systems, say to speed and simplify the operation of the combined system.

Consider two analogue digital systems AD^1 and AD^2 , with respective set of modes indexed by \mathcal{M}^1 and \mathcal{M}^2 . For ease of explanation, suppose they are running on different machines. We suppose that these systems are in general unrelated, so that the set of modes for the combined system \mathcal{M} *contains* the product of the modes for the subsystems $\mathcal{M}^1 \times \mathcal{M}^2$. If there is no cooperation between AD^1 and AD^2 , we will simply have $\mathcal{M}^1 \times \mathcal{M}^2$, with \mathcal{M}^1 being executed in parallel to \mathcal{M}^2 .

What would a cooperation between systems AD^1 and AD^2 look like? Suppose that the designers have identified modes $\alpha^1 \in \mathcal{M}^1$ and $\theta^2 \in \mathcal{M}^2$ as a possible cooperation. The default is that $\alpha^1 \in \mathcal{M}^1$ and $\theta^2 \in \mathcal{M}^2$ will still be executed in parallel.

However, if we are in $(\alpha^1, \theta^2) \in \mathcal{M}^1 \times \mathcal{M}^2$, an attempt will be made to start a cooperative mode. This new mode might consist simply of modified versions of $\alpha^1 \in \mathcal{M}^1$ and $\theta^2 \in \mathcal{M}^2$ which were optimised to run together, or at the other extreme it might consist of a completely new single process taking the place of both $\alpha^1 \in \mathcal{M}^1$ and $\theta^2 \in \mathcal{M}^2$ and running on a single machine.

11.3 A multi-scale hierarchy of modes

What roles do decomposition and modularity, encapsulation and inheritance, play in implementing multi mode analogue-digital systems? The behaviour of the physical system, and therefore the portfolio of models, has plenty of structure. Consider our autonomous car race in subsection 3.1:

Suppose the team decided that it would improve performance and safety to take account of the track conditions, specifically, if the racing track was dry or wet. As the code for mode α (the car is on the track and some distance from the nearest cars) was already quite complicated, they did not add separate cases to it. Instead they split α into two sub-modes, α -dry and α -wet. The picture $\text{State}_{\alpha\text{-dry}} = \text{State}_{\alpha}$ and $\text{State}_{\alpha\text{-wet}}$ was state_{α} with an additional state variable for just how wet the track was. Now only part of the code for α needed to be rewritten to give the codes for α -dry and α -wet.

Further, they decided that the old mode evaluation functions P_{α} were working perfectly well, so they did not rewrite them. Only if mode α was chosen did a further choice have to be made, between α -dry and α -wet. They also needed to decide if mode α was to be an *active* mode, i.e., a mode that could control the system, or a purely *passive* mode whose only purpose was to choose between the sub-modes α -dry and α -wet and to calculate the evaluation functions. They

decided that mode α should be active, as it was a perfectly good compromise solution if for some reason neither α -dry nor α -wet could be chosen reliably.

The answer to the question of decomposition comes from topology. A *refinement* of a given open cover for a topological space is an open cover – i.e., collection of open subsets whose union is the whole space – and so that any subset is completely contained in some subset of the original *coarser* cover.

A refinement is like increasing the resolution in a picture: we will resolve modes into submodes, an idea to be taken geometrically literally if we compare the nerve of the cover and the nerve of its refinement.

The set of modes \mathcal{M}' is a *refinement* of \mathcal{M} if for every $\xi' \in \mathcal{M}'$ there is an $\alpha \in \mathcal{M}$ so that ξ' is contained in α . We take this containment to be both physical (the subset of the physical state space corresponding to ξ' is a subset of that corresponding to α), and virtual (the software for ξ' is inherited from that for α).

In our racing example above, the refinement \mathcal{M}' of the original modes \mathcal{M} of Section 3.1 contains α -dry and α -wet, and we have the containments α -dry $\subset \alpha$ and α -wet $\subset \alpha$.

One complication that arises is that of multiple inheritance (see e.g., [33]). If the $\xi' \in \mathcal{M}'$ above is also contained in $\beta \in \mathcal{M}$, then we also need ξ' to inherit from β .

In fact, what we have been describing resembles a generalisation of the manifold data type, the idea of pre-sheaf [13], where data is localised to subsets of a set. Given a collection of subsets of a set, a *presheaf* assigns a structure $\mathcal{P}(U)$ to every U in the collection, together with a map (called the restriction map) $\phi_{VU} : \mathcal{P}(U) \rightarrow \mathcal{P}(V)$ for every $V \subset U$ in the collection.¹⁵ The basic idea is quite simple, given information about a set, we can restrict it to information about a subset.

The idea of refinement gives an implementation of the idea of presheaf, in that we have a restriction map from \mathbf{state}_α to every $\mathbf{State}_{\xi'}$ for every $\xi' \in \mathcal{M}'$ with $\xi' \subset \alpha$. Such a restriction map is used in going to finer scale structure in a multi-scale hierarchy. In our case we would also have to be able to move to coarser scales, which is not usually implemented in the structure of a presheaf. However one definition in sheaf theory might be useful: the idea of reconstructing data on a coarser resolution by combining finer resolutions. A sheaf is a special case of a presheaf where, given subsets $V_i \subset U$ where the union of the V_i is all of U , then we can reconstruct $\mathcal{P}(U)$ from the collection of the $\mathcal{P}(V_i)$.

11.4 Security

One problem with distributed systems is verifying that other parties are who they say they are. We can incorporate security in the data communications

¹⁵ There are several generalisations of presheaves, which eventually lead to just a functor between categories.

without greatly increasing the complexity of the software.¹⁶ As the only part of the system directly connected to external devices, the interface is the obvious place to add security. As the external devices are physically changed or have their software upgraded, so the security software associated with the interface must be altered.




Figure 14: Architecture of the interface with security

The idea is to *keep the details of the coding for security and error correction away from the control systems*. However, the control system has to know when its data may be corrupted. A limited number of flags may be passed from the interface to State_α describing problems with the data. For example, we might have:

Normal: Expected behaviour.

Link down: The device is not answering.

Noise: More errors in data transmission than normal, probably corrected.

Corrupted: Errors in data transmission which may not be corrected.

Authentication error: The device has not authenticated itself properly.

External interference: The security of the device may be compromised.

Cyber attack: Possibility that other devices may also be compromised, even if they verify as secure.

References

1. Ambaram, T, Beggs E J, Costa J F, Poas D, Tucker, J V, An Analogue-digital Model of Computation: Turing Machines with Physical Oracles, in: A. Adamatzky (ed.), Advances in Unconventional Computing, Emergence, Complexity and Computation 22, Springer-Verlag, 2016, 73-115. ,

¹⁶ There might be unintentional wireless interference between systems. There may be an unauthorised physical connection (e.g. an unapproved sensor in an autonomous vehicle). There may be a deliberate cyber attack.

2. Adcroft A, Campin J-M, Hill C & Marshall J, Implementation of an Atmosphere-Ocean General Circulation Model on the Expanded Spherical Cube, *Monthly Weather Review*, 132 (2004), 2845-2863.
3. Alexandroff P, Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung, *Mathematische Annalen* 98 (1928), 617-635.
4. Arnold, V I, *Ordinary Differential Equations*, The MIT Press, 1978.
5. Arnold, V I, *Mathematical Methods of Classical Mechanics*, Springer-Verlag, 1989.
6. Beggs, E J , Costa, J F, Loff B and Tucker, J V, Computational complexity with experiments as oracles, *Proceedings of the Royal Society Series A*, 464 (2008) 2777-2801.
7. Beggs, E J , Costa, J F, Loff B and Tucker, J V, Computational complexity with experiments as oracles. II. Upper bounds, *Proceedings of the Royal Society Series A*, 465 (2009), 1453-1465.
8. Beggs, E J , Costa, J F and Tucker, J V, Physical oracles: the Turing machine and the Wheatstone bridge, *Studia Logica*, 95 (2010) 279-300.
9. Beggs, E J , Costa, J F and Tucker, J V, Axiomatising physical experiments as oracles to algorithms, *Philosophical Transactions of the Royal Society A*, 28 (2012) 370, 3359-3384.
10. Beggs, E J , Costa, J F and Tucker, J V, The impact of models of a physical oracle on computational power, *Mathematical Structures in Computer Science*, 22 (2012) 853-879.
11. Beggs, E J, Costa, J F, Pocas, D, and Tucker, J V An analog-digital Church-Turing Thesis, *International Journal for the Foundations of Computer Science*, 25 (2014) 373-389.
12. Beggs, E J , Costa, J F and Tucker, J V, Three forms of physical measurement and their computability, *Reviews of Symbolic Logic*, 7:4 (2014), 618-646.
13. Bredon G E, *Sheaf theory*, McGraw-Hill, 1967.
14. Chen, M, Clayton, R, Holden, A V and Tucker, J V, Visualising Cardiac Anatomy Using Constructive Volume Geometry, Springer Lecture Notes in Computer Science, Volume 2674, Springer-Verlag, 2003, 30-38.
15. *Securing the future of German manufacturing industry: Recommendations for implementing the strategic initiative INDUSTRIE 4.0*, acatech-German National Academy of Science and Engineering, April 2013.
16. Einstein A, The foundation of the general theory of relativity, *Annalen der Physik* 49 (7) (1916), 769-822.
17. Henzinger T A, The Theory of Hybrid Automata, in M. Kemal Inan, Robert P. Kurshan (editors), *Verification of Digital and Hybrid Systems*, Volume 170 NATO ASI Series, 265-292, 2000
18. Hickey T J & Wittenberg D.K., Rigorous Modeling of Hybrid Systems Using Interval Arithmetic Constraints, in *Hybrid Systems: Computation and Control*, Springer Lecture Notes in Computer Science, Volume 2993, 2004, 402-416.
19. Hurevich, W & Wallman, G *Dimension theory*, Princeton University Press, 1948.
20. De Kleer J & Brown J S, A qualitative physics based on confluences, *Artificial Intelligence*, Special volume on qualitative reasoning about physical systems, Vol. 24 no. 1-3, (1984), 7-83,
21. Laskar J, Large-scale chaos in the solar system, *Astronomy and Astrophysics*, 287 (1) (1994), p L9-L12.

22. G. Labinaz and M. Guay, *Viability of Hybrid Systems*, Springer, 2012
23. Lee, E A, Computing needs time, *Communications of the ACM*, 52 (2009), 70-79.
24. Lee, E A and Tripakis S, Modal Models in Ptolemy, in 3rd International Workshop on Equation-Based Object-Oriented Modeling Languages and Tools (EOOLT). 2010. Oslo, Norway: Linköping University Electronic Press, Linköping University.
25. Lunze, J, & Lamnabhi-Lagarrigue, F, *Handbook of Hybrid Systems Control Theory, Tools, Applications*, Cambridge University Press, 2009.
26. Oded Maler, Amir Pnueli and the dawn of hybrid systems, *Proceedings of the 13th ACM international conference on hybrid systems: computation and control* ACM, 2010, 293-295
27. Misner C W, Thorne K S & Wheeler J A, *Gravitation*, W.H. Freeman, 1973.
28. Poole, M J, Tucker, J V and Holden, A V, Hierarchies of spatially extended systems and synchronous concurrent algorithms, in B M?ller and J V Tucker (editors), *Prospects for hardware foundations*, Springer Lecture Notes in Computer Science, Vol 1546, 1998, Springer-Verlag, 184-235.
29. Poole, M J, Tucker, J V and Holden, A V, Hierarchical reconstructions of cardiac tissue, *Chaos, Solitons and Fractals*, 13 (2002) 1581-1612.
30. Riemann, B, On the hypotheses which lie at the bases of Geometry, *Nature*, Vol. VIII (1873) Nos. 183, 184, pp. 14-17, 36, 37. (Trans. W K Clifford)
31. Rindler W, *Essential relativity: special, general, and cosmological*, 2nd edition, Springer, 1997.
32. Rudin, W, *Real and complex analysis*, 3rd edition, McGraw-Hill, 1987.
33. Schärli N, Ducasse S, Nierstrasz O & Black A, Traits: Composable Units of Behaviour, *ECOOP 2003 - Object-Oriented Programming*, Springer Lecture Notes in Computer Science, Volume 2743, 2003, 248-274
34. Segal G, Classifying spaces and spectral sequences, *Inst. Hautes Études Sci. Publ. Math.* No. 34 (1968), 105-112.
35. Spanier E H, *Algebraic topology*, Springer 1989.
36. Hirsch, M W, Smale, S, Devaney, R L, *Differential Equations, Dynamical Systems and an Introduction to Chaos*, Elsevier, 2004.
37. Thompson, B C, Tucker J V and Zucker, J I, Unifying computers and dynamical systems using the theory of synchronous concurrent algorithms, *Applied Mathematics and Computation*, 215 (2009), 1386-1403.
38. Truitt T D, Hybrid computation... What is it?... Who needs it?... *IEEE proceedings - Spring Joint Computer Conference*, 1964
39. Tucker J V and Zucker, J I, Computability of analog networks *Theoretical Computer Science*, 371:1-2 (2007) 115-146.
40. J V Tucker and J I Zucker, Computability of operators on continuous and discrete time streams, *Computability*, 3 (2014) 9-44.
41. Weihrauch, K, *Computable analysis*, Springer, 2000.